Article published in ITProfessional,
14(3):13,21, May-June 2012
doi: 10.1109/MITP.2012.25

NLM People Locator

A system for family reunification

George Thoma, Sameer Antani, Michael Gill, Glenn Pearson, Leif Neve
Lister Hill National Center for Biomedical Communications
US National Library of Medicine, Bethesda, Maryland

Abstract

One of the disturbing consequences of a natural or manmade disaster is that people go missing, with distraught family and friends having no information as to their whereabouts. Missing people often turn up in refugee camps or in hospitals where information on them may be taken and stored in a database. To use this information for family reunification, the Lister Hill Center, an R&D division of the National Library of Medicine, has developed *People Locator (PL)*, a Web site to which photos and metadata (name, age) for such missing (or found) people can be posted by hospital staff, relief workers, or family members, and which can be searched by professional counselors or the public. The Web site is designed to receive information through the Web, a hospital triage-specific application (TriagePic), an iPhone app (ReUnite), and through interfaces to other sites. This article describes the PL system and its components.

Introduction

In recent years the U.S. National Library of Medicine, as an information provider and technology developer, has sought to mitigate the effects of mass casualty events, both natural and deliberate, in America and abroad. Programs have therefore been initiated to provide relevant information to first responders and the public[1]. In addition, research and development projects are underway to design and deploy technologies to track hospitalized disaster victims, to aid triage operations, and to reunite families.

People displaced or injured during a disaster can lose contact with family and friends, but missing survivors often turn up in evacuation shelters or in hospitals. There, information on them may be taken and made available for searching by loved ones seeking them. For this purpose, the Lister Hill Center, an R&D division of the National Library of Medicine, has developed ***People Locator (PL)***, a Web site to which name, photos, and other information about a lost or found person can be entered by hospital personnel, relief workers, or family members, and which can be searched by professional counselors or the public.

People Locator was first deployed during the January 2010 Haiti Earthquake, and since then for several disasters including the Japan Tsunami and the New Zealand quake. A large number of records (photos

[1] Particularly through NLM's Disaster Information Management Research Center. http://disaster.nlm.nih.gov

and descriptions for missing people) were gathered by PL, some sent directly to this site, and many others imported from other similar systems, notably Google's Person Finder system.

The PL system

The principal components of People Locator are shown in Figure 1. At its core, a database holds data on missing and found people, and a search engine indexes this for efficient searching. There are multiple ways to enter data in the database, means to share data with other repositories with a similar goal, and a user interface offering multiple views of the data for browsing and searching.

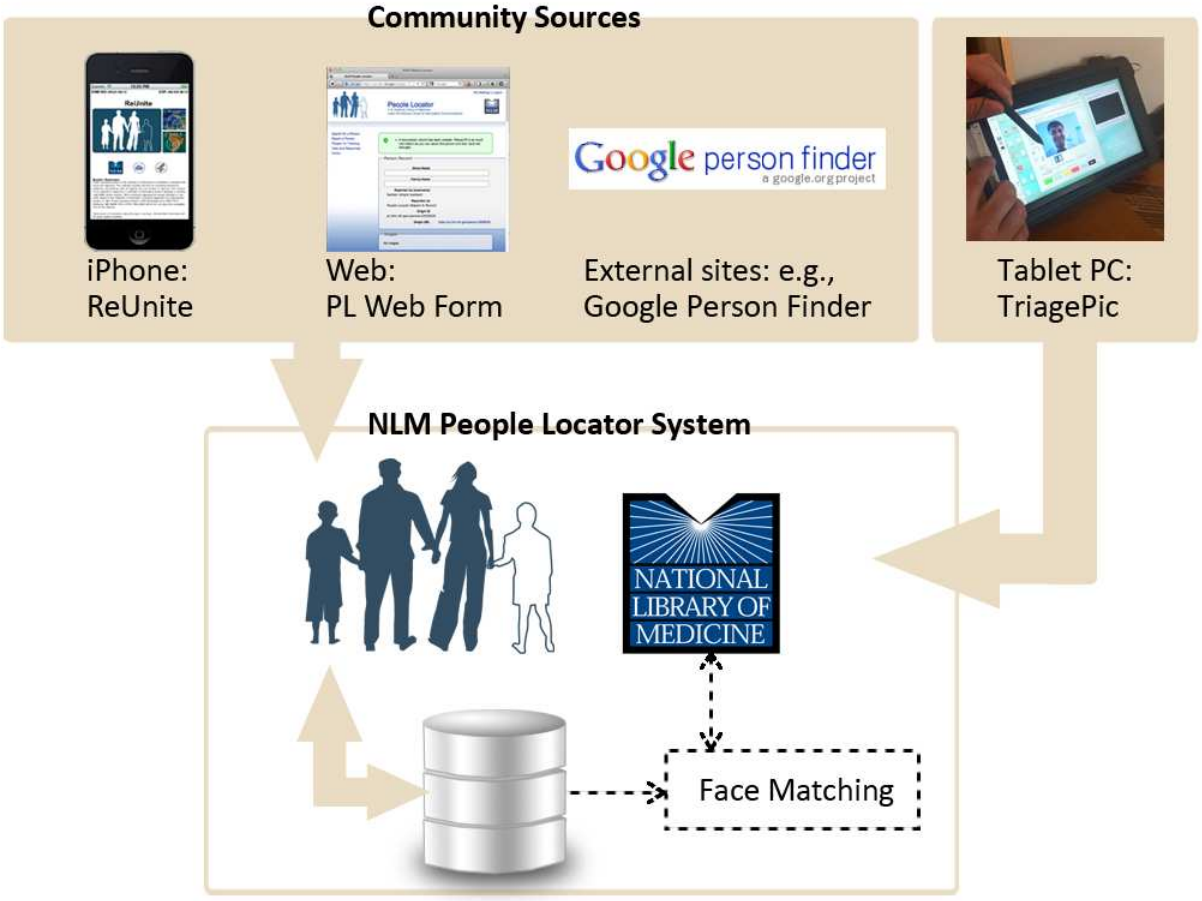

Figure 1: An overview of PL and its data input sources, described in the main text. R&D efforts seek to improve database search, including using a combination of text and image (face) based matching.

Broadly, PL supports two "flavors" of use: for either *hospital triage-based* or *community-wide* operation. The first aligns with triage procedures and workflow at a hospital for both data input and searching, while the community-wide system allows the uploading of data from anywhere in a large affected area (such as Haiti and its diasporas in the earthquake's aftermath), and public search.

For reporting missing and found people, PL offers multiple methods. The three most general (but particularly of importance to community-based outreach) are:

- through a (structured) PL Web form.

- through a structured form within ReUnite, an iPhone app.
- as unstructured text in the subject line of an email message (not shown in figure).

In addition, data can be imported from other repositories with which PL is interoperable, e.g., Google's Person Finder. For data gathering at hospitals, NLM offers TriagePic, a Windows-based application for laptop and tablet PCs.

At the PL Web search page, entering a missing person's name yields the corresponding photos and descriptive metadata records from the database. There is also a viewing feature, a "notification wall", to dynamically display all those reported missing and found, or clustered by known status (alive and well, injured, or deceased), gender, and age categories.

Database design

The initial design of the PL database and Web site was based on the circa-2008 Sahana disaster management system, developed after the Indian Ocean tsunami [1][2]. Following the approach adopted by the Sahana developers, PL employs a relational database design centered on "person records" stored in the *person_uuid* table. Every report of a person generates two entries in this table, one for the reporter and another for the reported individual. Many of the other tables in the system are child tables of *person_uuid*, recording some aspect of these individuals. For example, *person_status* contains status about the health or location of reported persons, and *person_physical* records their physical characteristics. The *contact* table saves contact information for both reporters and reported persons. Information on reporters allows reported persons to contact those looking for them, and contact information for the missing makes it possible for reporters to closely track people they are looking for.

Other important tables in the system are *hospital*, for associating reported persons with hospitals, and *incident*, for associating them with a specific disaster event. The *user* table contains information on all registered users of PL, including those who have reported a missing or found person. The *image* table contains information on photos, such as URL and pointers to where they reside in local storage.

As PL evolved beyond the original Sahana functionality, it introduced new tables. For instance, one set is for sharing data imported from other family reunification sites via the PFIF standard (explained below). The *pfif_person* and *pfif_note* tables capture records (explained later) from other repositories, and the *pfif_repository* table lists known PFIF repositories, that is, those that are interoperable with PL. The remaining PFIF tables control the automated export and import processes.

Among other new features, records in PL are indexed and searched for by Solr, a Java search platform from the Apache Lucene project[2]. Solr offers better performance and more powerful search string relevance matching than do SQL queries. On the downside, it introduces a small delay between the time at which persons are reported and when they can be queried, since Solr needs to first index new incoming data.

[2] http://lucene.apache.org/solr/

Sharing data with other repositories

To promote family reunification on a wide scale, PL is designed to exchange missing person data with other survivor repositories. As mentioned above, we share data using the *People Finder Interchange Format* (PFIF), an open XML-based standard for exchanging and aggregating missing and displaced person data.[3] It was created in September 2005 in the aftermath of Hurricane Katrina, and adopted by Google, Yahoo, and others to share missing person data. It saw extensive use again after the Haiti earthquake in January 2010 when Google launched their Person Finder and exchanged PFIF data with our system, as well as with CNN and the New York Times.

PFIF data includes *person records* containing identifying information about a person, and *note records* which contain comments and updates on the status and location of a person. A typical missing person description consists of a single person record and multiple note records. PFIF has matured since its first use in Katrina to better accommodate the needs of international users regarding postal addresses, personal names, and Unicode. It also includes a mechanism to eliminate ("expire") records, if appropriate.

Our system is equipped to exchange missing person data with any PFIF-compliant repository by means of regular automated exports and imports. In the aftermath of the Christchurch Earthquake in February 2011 and Tsunami/Earthquake in Japan in March 2011 [3], we implemented this automated data sharing with Google Person Finder. As a result, missing person data reported on either Web site appeared seconds later at the other, allowing users around the world to report and search on the same data at the site they preferred.

We used the Google Person Finder data API[4] to undertake this data sharing. We have subsequently implemented part of that API in our system to allow other repositories direct access to our data, particularly important when PL might be the only repository, or the first one, deployed in response to a disaster.

User interface

Arriving at the PL Web site, a user first selects a particular disaster event, such as "Haiti Earthquake", and reports on a missing or found person, or searches the database. Searching for or reporting a missing person is then done for that particular event, not across all events. Once the event is selected, the user is presented a text search box. Clicking the search button yields all records collected for the event (Figure 2). No login is required to search for missing people, however registration is required for reporting data on people (explained below). The public may search for a missing person by (a) a name, or partial name using a wildcard (for example: "Cath*" will find "Catherine"), (b) using the word "unknown" to search records lacking names, or (c) a blank box, which delivers all records in the repository.

[3] http://zesty.ca/pfif/

[4] http://code.google.com/p/googlepersonfinder/wiki/DataAPI

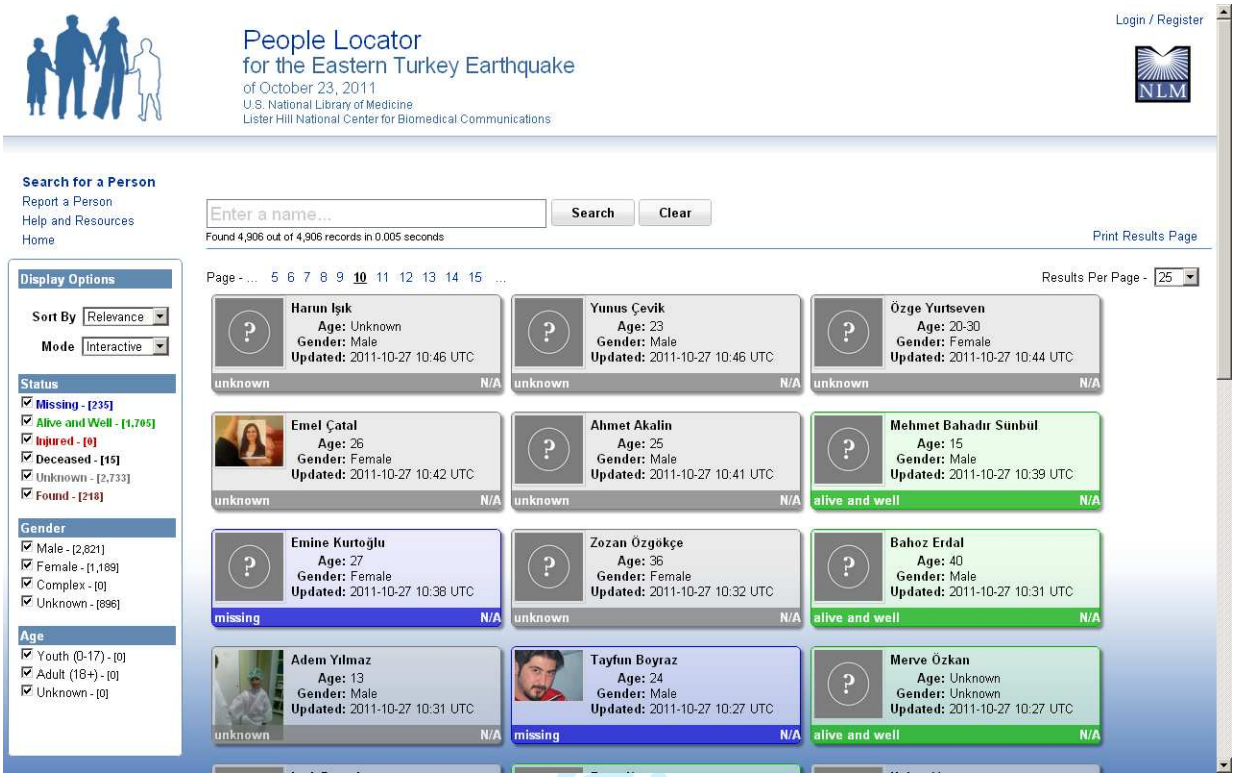

Figure 2: PL search interface (Interactive Mode) showing missing person records from the Eastern Turkey Earthquake of October 2011.

Search results may be displayed in three user-selectable modes: (a) Interactive (the default), (b) Hands-Free, and (c) Full-Screen or "Notification Wall". In all three modes, one may filter the results using checkboxes for gender, age and status (alive and well, injured, etc). The user may also re-order the displayed data according to the time the record was posted or updated, name, age or status. Clicking on a retrieved record yields a more complete (but still partial) record, and clicking again displays the full record stored in our system. The user may also print out these search results.

While no login is required for searching, registration is required if the searcher wants to be notified of any future matches or status updates. This follows from the assumption that before one reports a missing person, he/she is likely to first check whether a report already exists in our system. Registration is a simple process requiring the user to provide their first and last name, desired user name (login ID), email address and a strong password.

The Interactive mode shown in Figure 2 displays results in a row and column layout with a thumbnail photo, as well as name, age, gender, status and the last-updated timestamp. The Hands-Free choice arranges records as a film-strip that automatically scrolls from left to right (Figure 3). The Full-Screen or Notification Wall mode does the same thing, but the results occupy the entire screen, better for display to a larger audience -- in an auditorium, for example. These latter two modes are more suitable when the user needs to review many pictures at once, and were designed for Emergency Management or

counseling staff to rapidly and conveniently view records of incoming patients. Video player-like controls provide the flexibility to start, pause, rewind, and slow down or speed up scrolling.

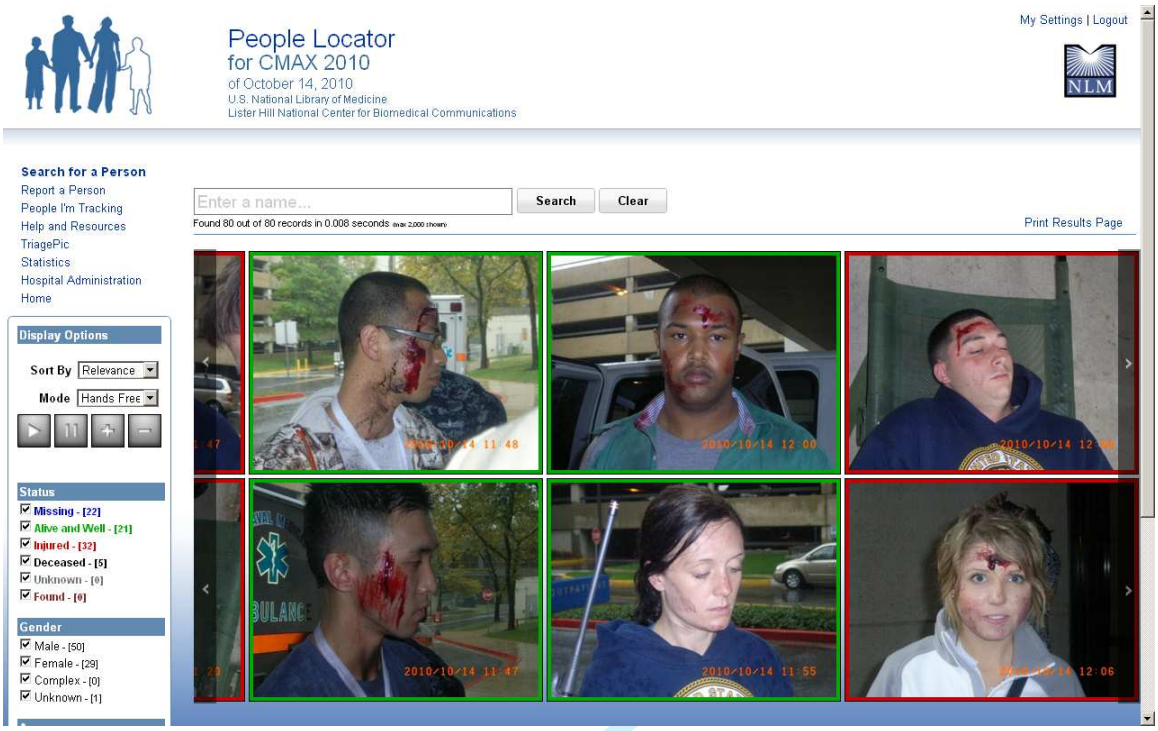

Figure 3: Hands-Free mode showing records from the October 2010 CMAX Drill (a disaster event exercise held jointly by three neighboring hospitals in Bethesda, Maryland)

Ways to report missing or found people

As mentioned earlier, there are multiple ways to report data on people. These have been broadly classified into two categories: (i) methods suitable for community reporting; and (ii) more specific to a hospital scenario. Here we provide more detail on the form on the PL Web site and ReUnite that are suitable for use anywhere, and the TriagePic application that is specific to hospital operations.

PL Web site – for community reporting

In addition to searching the PL also has a Web-based form for reporting data on people. This action requires registration.

As shown in Figure 4, registered users can select the relevant disaster event and provide the requested information such as the person's name, age (or age range) in years, gender, eye color, skin color, height, weight, and any other distinctive features. Last seen or last known location may be provided in descriptive terms. Private contact information can also be provided, but while this is not visible to the public, it is available to the site administrators to aid in family reunification.

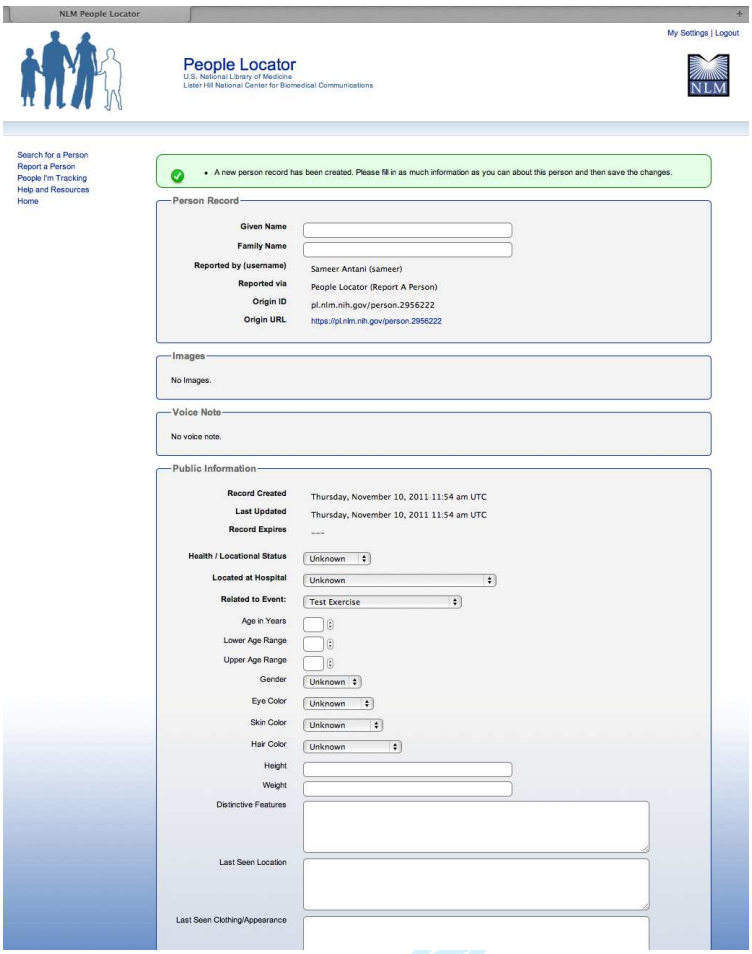

Figure 4: The PL Web form to report data on people.

ReUnite, an iPhone app – for community reporting

As recent events have demonstrated, there is active community participation in the recovery and reunification efforts after a disaster. This was evident from those seeking the missing after the 9/11 terrorist attacks on the World Trade Center in New York City, and more recently after the Haiti Earthquake of 2010, and the Japan Earthquake of 2011. In each instance people put up signs of those missing on community bulletin boards and other media. Since the PL system provides an electronic bulletin board where the missing can be reported, it needs a wide variety of easy-to-use mechanisms for the data to be provided to it. One such method is the ubiquitous smartphone.

Following the Haiti Earthquake, in a "tiger team" approach, we developed an iPhone app called "ReUnite" (initially called "Found in Haiti") that is also compatible with the iPod Touch and iPad devices. The app provides a structured form for reporting missing and found people. Developing for the iPhone offered several advantages: (i) ubiquity, (ii) broad applicability (iPhones, iPodTouch, iPad), (ii) standardized hardware, (iii) powerful software development platform, and (iii) a uniform distribution mechanism. Some of these advantages may also be found with other mobile-software development platforms for widely-available devices (e.g., Android, Blackberry, Win Phone).

Figure 5: Screen shots of ReUnite iPhone app: (a) Initial screen; (b) Form for recording person information; (c) Image capture and tagging capability; (d) Summary information; (e) and (f) Mapping feature that uses GPS to identify geographic and postal information.

From its inception, the app design focused on two key capabilities: (i) report missing and found people; and (ii) search PL. Also, it was designed for two kinds of users: (i) the layperson; and (ii) professionals such as a social worker or relief worker. A mother looking for her missing family member modeled the layperson. She would use ReUnite to provide the missing person's name, age or range, gender, a picture, location where the person was last known to be, and contact information where the person's status could be reported. A social worker, on the other hand, would typically be at a recovery camp where she would be reporting those that are found. In this case, the location is known, as is the person's health status, along with a picture and perhaps a voice note that could comfort those looking for the individual.

The ReUnite app, which was highlighted by Apple iTunes Store as one of the "New and Noteworthy" apps, combines all of these desirable characteristics. As shown in Figure 5, it provides a form view in which the name, gender, age, health status, photo (with tags/annotations), and location (aided by GPS and Google Maps) can be recorded. Also the text and voice comments fields provide the capability to add helpful information.

The user can also search for records from within the app by using the embedded browser. Future enhancements to the app include improved search and filtering capability, native within-app search, and enhanced interaction with the PL Web site.

TriagePic – for hospital-based reporting

This Windows-based application is designed to quickly gather photos and minimal information about disaster victims arriving at a hospital, and to forward this data for display on the PL Web site.

In its first version, photos were taken by a hospital person with a special Bluetooth camera paired with a laptop hosting TriagePic. When within a dozen feet of the laptop, the camera operator could send a photo to TriagePic, where related text data was entered, usually by another staff member. During drill testing, this proved cumbersome and error-prone, requiring too much synchronization and verbal relaying between photographer and data entry person.

Subsequently, these functions were combined in a tablet computer (or laptop with webcam), so that a single staff member can:

- take the photo
- enter a mass casualty ID
- record gender, whether adult or child, and a name if time allows
- route the patient to a hospital zone for treatment (color coded as green, yellow, red, etc. depending on severity of injury).

This information is immediately sent to PL by Web Service, and optionally to email recipients.

In a recent region-wide exercise, Capital Shield 2012, we successfully fielded two Motion Computing CL900 tablet computers with TriagePic at Suburban Hospital (Figure 6). Each tablet runs Microsoft Windows 7 Professional, is finger-touch sensitive, and has a digitizer pen and two built-in cameras. This proved considerably more successful than the first approach. The collected data was transmitted to PL through the hospital's enterprise WiFi network. Each photo is jpeg-compressed, speeding up transmission by WiFi while retaining good quality. Compression is a frequent requirement of wireless telemedicine applications (e.g., Olariu et al [4]).

TriagePic works best with this tablet in a "landscape" orientation. To use the rear-facing 3-Megapixel camera, the tablet is carried, not docked. Hospital staff liked the tablet's optional holder, which has elastic hand straps on the back and a shoulder strap. If the tablet were docked, a physical keyboard would be the best mode of text entry. Instead, a Windows 7 virtual keyboard and/or handwriting recognition is used, for which the digitizer pen is helpful. TriagePic is designed to exploit the finger-touch feature in the tablets for non-text entry, for example, via enlarged checkboxes to specify gender and adult/child.

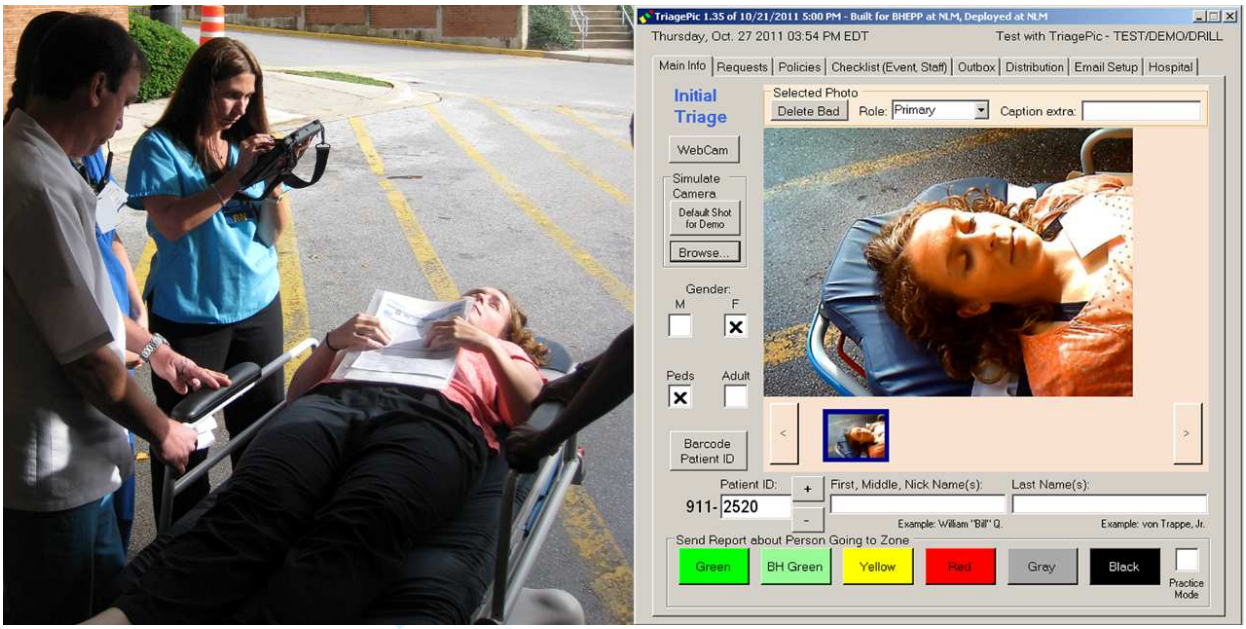

Figure 6: TriagePic in use on a tablet computer during the Capital Shield 2012 Drill at Suburban Hospital-Johns Hopkins Medicine in Bethesda, Maryland. On the right, the TriagePic interface.

Future steps

Challenges remain in several areas. For retrieval, we intend to explore search by face matching (discussed next) and multimodal (text + image) search strategies (e.g., Rojas [5]). Also discussed below: strategies to notify the public and relevant government and emergency response organizations through SMS text messaging and social media outlets (Twitter, Facebook, Google+). We will advance interoperability with other disaster mitigation sites through common data formats and data retention policies.

There are opportunities for hospital-focused improvements. For example, the system could be considered for routine ER use, not just disaster response, if it had privacy and security enhancements, such as HIPAA-compliance [6]. It would also benefit from integration with other hospital systems, such as real-time patient trackers using RFID/IR tags [7].

Face matching and photo de-duplication

We want to extend the search technique to *image queries*. That is, a family member should be able to send in a photo of a loved one, and get the top (say, 5) most similar pictures in the database. This face matching requires efficient techniques for the extraction of image features (shape, color, texture) and also annotated validated (ground truth) datasets. Ongoing research has resulted in a prototype for face location and identification using tools in the OpenCV[5] image processing library. Faces are located using the method proposed by Viola and Jones [8], and face descriptors computed using SURF[6] [9] (*speeded up*

[5] OpenCV: Open Source Computer Vision library: http://opencv.willowgarage.com/wiki/

[6] SURF: Speeded Up Robust Features: http://en.wikipedia.org/wiki/SURF

robust feature). Localized faces were indexed by computing these SURF descriptors, and those within a threshold similarity score were considered similar.

A related problem concerns the same photo appearing multiple times in the database, possibly because several people sent it in during a disaster event. Eliminating these duplicates would reduce the search space and thereby improve performance. Several methods were investigated, and a C++ prototype was developed to index the image collection based on the Haar Wavelet[7] image descriptors, retaining positions and signs of only the most significant coefficients. Measuring the matching distance between photos and grouping those below a threshold distance gives sets of near-duplicates.

Challenges in this area include low-resolution photos, typical of those from cell phones, faces that are occluded with clothing or hair, and variations in pose and ambient lighting.

Notification and data exchange strategies

The use of social media in disasters for notifications, reunification, help requests, situational awareness, and general emergency management is becoming increasingly important [10][11].

We therefore seek to notify as many interested parties as possible through the use of multiple channels when we deploy PL for a disaster event. In order to do so efficiently, we plan to build in automatic notification via social networking sites, (Twitter, Facebook, etc.), SMS, the Integrated Public Alert and Warning System (IPAWS[8]), LISTSERVs, email lists and other avenues. Standardized message creation and distribution strategies have to be developed and integrated into the system. Evaluation of distribution strategies has to be accomplished in order to judge their effectiveness.

There are a number of potential partners for disaster information exchange. In the U.S., the American Red Cross provides a "Safe and Well" Web site, and coordinates with other entities through a shelter network. To complement the National Center for Missing and Exploited Children, the U.S. Federal Emergency Management Agency (FEMA) attempted a National Emergency Family Registry and Locator System (NEFRLS)[9] for adults. Information about missing people brought to police attention is disseminated through the FBI's NCIC 2000 system. Traditional patient transport/tracking systems include the military's TRAC2ES[10] and DHHS's JPATS[11]. We are particularly interested in the emerging Tracking Emergency Patients/Clients protocols [12] supported by FEMA and U.S. Department of Health and Human Services (DHHS), which should allow exchange of messages among emergency managers at the local, state, and federal level, using both IPAWS-OPEN and private networks.

[7] Haar Wavelet: http://en.wikipedia.org/wiki/Haar_wavelet

[8] IPAWS: http://www.fema.gov/emergency/ipaws

[9] NEFRLS: https://egateway.fema.gov/inter/nefrils/home.htm

[10] TRAC2ES (TRANSCOM Regulating and Command & Control Evacuation System): https://www.trac2es.transcom.mil

[11] JPATS (Joint Patient Assessment & Tracking System): http://teams.hhs.gov/jpats

The US State Department is involved in international responses, as is the UN Office for the Coordination of Humanitarian Affairs (OCHA). Important sites include the International Committee of the Red Cross's Family Links, and Missing.Net of Foundation Casques Rouges.

Acknowledgments

This work is supported by the Intramural Research Program of the National Institutes of Health, the National Library of Medicine and the Lister Hill National Center for Biomedical Communications. The authors acknowledge the significant contributions of Greg Miernicki, Merwan Rodriguez, and Michael Chung in various aspects of system design, development and deployment of People Locator, contributions by Eugene Borovikov and Girish Lingappa to ongoing research in face matching, and Bathiya Senevirathna and Krittach Phichaphop for development of the ReUnite iPhone app. Finally, we acknowledge the collaboration of the Sahana Software Foundation[12] and institutions in the Bethesda Hospitals Emergency Preparedness Partnership (BHEPP) [13].

References

1. Careem M, De Silva C, De Silva R, Raschid L, Weerawarana S. Sahana: Overview of a disaster management system. Intl Conf Info and Automation (ICIA 2006), Dec 2006. ISBN 1-4244-0555-6. DOI 10.1109/ICINFA.2006.374152 .
2. Samaraweera I, Corera S. Sahana victim registries: Effectively track disaster victims. Proc. 4th Int Conf on Info Sys for Crisis Response and Management (ISCRAM '07), 2007.
3. Arifumi Utani, Teruhiro Mizumoto, Takashi Okumura. How geeks responded to a catastrophic disaster of a high-tech country: rapid development of counter–disaster systems for the great east Japan earthquake of March 2011. Proc Spec Workshop on Internet and Disasters (SWID '11), ACM, 2011. ISBN 978-1-4503-1044-4. DOI: 10.1145/2079360.2079369.
4. Olariu S, Maly K, Foudriat E, Yamany S. Wireless support for telemedicine in disaster management. Proc 10th Intl Conf Parallel and Distr Sys (ICPADS 2004), July 2004: 649-656.
5. Rojas J. Assessment of a Proprietary Online Smart-Family-Matching Tool to Reunite Lost Families. IEEE Africon 2011, (Livingston, Zambia, Sept 13-15):1-6, 2011. 978-1-61284-993-5/11/
6. Hodge J. Health Information Privacy and Public Health. J Law, Med & Ethics, 31. Am Soc of Law, Med & Ethics, 2003:663-671.
7. D'Souza I, Ma W, Notobartolo C. Real-Time Location Systems for Hospital Emergency Response. IT Professional, 13(2), IEEE Comp Soc, March-Apr 2011, 37-43. DOI 10.1109/MITP.2011.31
8. Viola P, Jones MJ. Robust Real-Time Face Detection. Int. J Computer Vision, 57(2):137-54. 2004. DOI: 10.1023/B:VISI.0000013087.49260.fb
9. Bay H, Ess A, Tuytelaars T, van Gool L. SURF: Speeded Up Robust Features. Computer Vision and Image Understanding, 110(3):346-59, 2008.

[12] Sahana Software Foundation: http://www.sahanasoftwarefoundation.org

[13] Bethesda Hospitals Emergency Preparedness Partnership (BHEPP): http://www.bhepp.org

10. Yates D, Paquette S. Emergency knowledge management and social media technologies: A case study of the 2010 Haitian earthquake, Intl J Info Man,31(1) Feb 2011,6-13. DOI: 10.1016/j.ijinfomgt.2010.10.001
11. Gao H, Wang X, Barbier G, Liu H. Promoting Coordination for Disaster Relief – from Crowdsourcing to Coordination. Proc 4th Intl Conf Social Computing, Behavioral-Cultural Modeling and Prediction (SBP '11). Springer-Verlag. ISBN 978-3-642-19655-3.
12. EDXL Practitioner Steering Group, TEP/TEC Steering Committees, sponsored by DHS. Draft specifications are available in TEP and TEC folders under www.evotecinc.com.